

# BITS

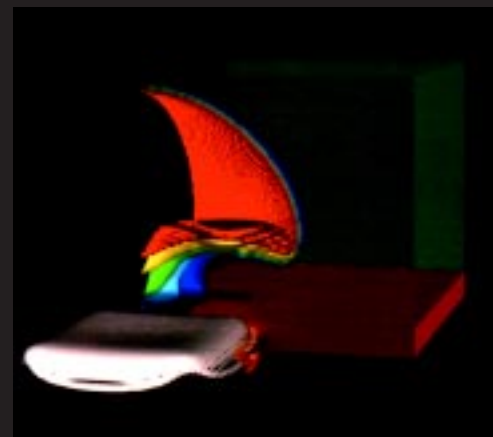
*computing & communications news*

---

## *Some of What's Inside*

**High Performance Computing**  
Using Tecolote Components to  
Extend Object-Oriented  
Programming

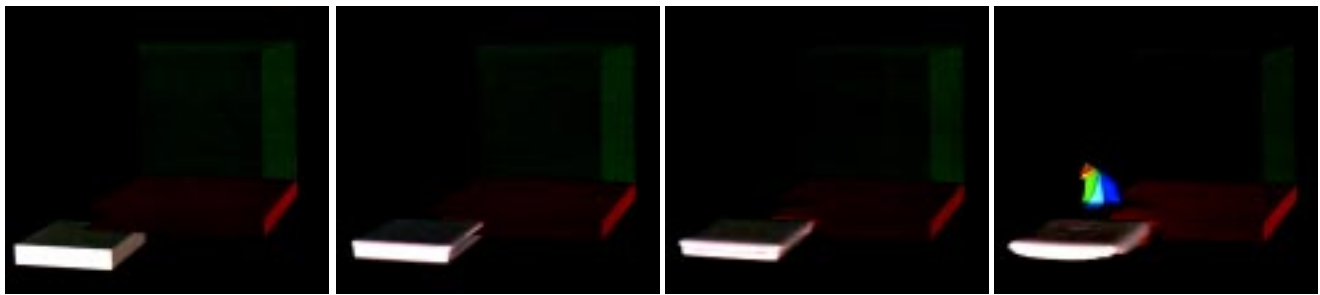
**Desktop**  
Desktop on Demand: A Desktop  
that Displays on Your  
Computer, but Actually Runs on  
the Enterprise Server



**Information Systems**  
Laboratory Overview Presentation  
Materials on the Web

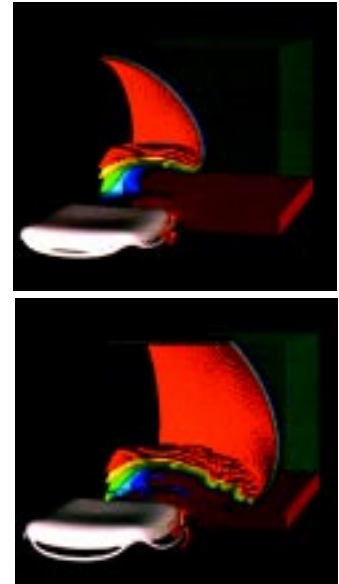
**What's Happening**  
New Electronic Resources from  
the Research Library

**Los Alamos**  
NATIONAL LABORATORY  
September 1999



The sequence of images result from a three-dimensional simulation that tests a reactive high-explosive burn model. The model is a component developed within the Tecolote Framework. The sequence shows the evolution of a shock-to-detonation transition from the impact of an aluminum fragment striking an aluminum plate covering a block of PBX-9501 high explosive (HE). After impact, a shock wave is generated that propagates into the HE causing it to react. The extent of reaction is shown by the colored isosurfaces with each isosurface representing a constant value of the HE burn fraction. Blue corresponds to a burn fraction of 0.0 (solid HE) while red corresponds with a burn fraction of 1.0 (gas products). Other colors correspond to levels of intermediate reaction. Simulation performed by Ed Idar, X-NH. For more information contact Idar at (505) 667-1721, or e-mail [esi@lanl.gov](mailto:esi@lanl.gov).

(Editor's note: BITS is featuring a series of articles on Tecolote, which began in August 1999 with the article *Tecolote: An Object-Oriented Framework for Hydrodynamics Physics*. In this issue see *Using Tecolote Components to Extend Object-Oriented Programming* in the High Performance Computing section.)



Produced by the Computing, Information, and Communications (CIC) Division

Design: Donald Montoya and Julie Medina (CIC-1)

Managing Editor: Denise Sessions, [denise@lanl.gov](mailto:denise@lanl.gov) (CIC-1)

Illustration: Dave Delano

Photography: John Flower (CIC-9)

Printing: Imaging Services Group (CIC-9)

**BITS Contributors' Board:** Hal Marshall, CIC-6; Kimberlyn Mousseau, CIC-15; Gina Fisk, CIC-5; Kathleen Jackson, CIC-3; Kei Davis, CIC-19; Nikki Gaedecke, CIC-6; Octavio Ramos, CIC-1/NMT-DO; David Van Etten, CIC-6, Denise Sessions, CIC-1/CIC-6, and Don Willerton, CIC-6 Group Leader.

**Attention Subscribers:** There are two ways you can update your subscription data. E-mail [bitsupdate@lanl.gov](mailto:bitsupdate@lanl.gov), or fill out the form inside this issue. Internal LANL subscribers please note that until we have a mechanism in place to verify our subscription database with the LANL Employee Information System we will rely on you to update your mailstop. Thanks for keeping us updated on subscription data.

# Table of Contents

<b>Information Systems .....</b>	<b>2</b>
Laboratory Overview Presentation Materials on the Web .....	2
 <b>Desktop Computing .....</b>	 <b>6</b>
Desktop on Demand: A Desktop That Displays on Your Computer, but Actually Runs on the Enterprise Server .....	6
Grassroots Software Management: Simple Things to Do .....	8
News About Microsoft Networking at Los Alamos: The LANL Master Accounts Domain Is Up and Running! .....	10
Notes on Setting Up and Configuring an NT Domain .....	12
 <b>Infrastructure .....</b>	 <b>15</b>
Web Content Architecture: Taming the Tangle of Protocols .....	15
 <b>High Performance Computing .....</b>	 <b>19</b>
Using Tecolote Components to Extend Object-Oriented Programming .....	19
 <b>What's Happening .....</b>	 <b>23</b>
CIC Division Hosts Student Poster Presentations .....	23
New Electronic Resources from the Research Library .....	26
Research Library Resources .....	27
Software Documentation Services .....	27
Technical and Advanced Computer Training Courses .....	28
BITS Subscription/Comment Form .....	29
1999 12-Month Index .....	31

## Laboratory Overview Presentation Materials on the Web

*by Denise Sessions, BITS Managing Editor, and Chris Lindberg, Web Designer Programmer, CIC-1 Communication Arts and Services*

If you're creating a presentation about Los Alamos National Laboratory, wouldn't it be nice to use professionally created viewgraphs about the Lab? In an effort to provide a source of consistent, up-to-date presentation information about the Laboratory, George Kwei with the Director's Office requested that CIC-1, Communication Arts and Services, create a common enterprise repository to capture viewgraphs of Lab corporate history accessible to everyone at the Lab. The solution turns out to be a Web-based document management hosting service.

Currently the repository is on a Web site called the "Lab Overview." Ed Borrego, graphic illustrator in the Director's Office, has created several hundred viewgraphs of which 80 reside on the Web site. Borrego is enthusiastic about the new tool that will automate version control and help manage slide inventory. As the graphic illustrator for this project, Borrego is responsible for evaluating and pulling together the numerous elements needed to convey consistent, accurate, and reliable information about the Lab for public consumption. Eventually the site will contain not only Laboratory overview viewgraphs, but there will be perhaps as many as

several thousand viewgraphs describing a much broader range of Lab activities. The Director's Office will be seeking help to add viewgraphs to the site. The Web site is on the Lab's Intranet in the blue, unclassified-protected side of the internal network at this URL: <http://www.lanl.gov/laboverview>.

## Web-Based Document and Content Management

CIC-1's Chris Lindberg took on the assignment to choose the appropriate software tool and build the Web site. At Kwei's request, Lindberg evaluated off-the-shelf software products designed to manage documents and content using an automated Web-based system. He chose products he thought would do the best job in meeting the specifications of this project.

## Project Specifications

The purpose of the project is to make the Laboratory Director's electronic presentations available to other presenters at the Laboratory. The tools that were evaluated had to meet a variety of technical specifications. Top consideration was given for cost-effectiveness and versatility of user options. The media should be available in Microsoft® PowerPoint® and Web-accessible formats, such as HTML and PDF. For accessibility the viewgraphs and photographs needed to be indexed and searchable by title, organization, topic, author, photographer, and program. Instead of several versions of a viewgraph floating around, only the latest version should be available.

Some control of the review process—or workflow—should be automatic so that the necessary reviews for quality and classification could be included. Another requirement was that only contributors whose names appear on an access-control list be able to submit presentations or resources to the site. Because of the number of contributors, revision control was a core requirement. The system needed to be compatible with archived materials.

## Evaluating Available Commercial Software Solutions

Lindberg compared Documentum® and IntraNet™ Solutions' "Intra.doc!" and found them to be very similar in many ways. Both products satisfy the project's core requirements.

Lindberg observed the versatility and customization capability of Documentum, noting that "given enough time and resources, virtually any functionality can be achieved." To create queries, Documentum requires writing standard/structured query language (SQL) statements and other programming. Intra.doc! generates the appropriate SQL statements easily through its GUI Web Layout Editor, which is a Java program that dynamically updates the layout and hierarchy of the Web pages. No knowledge of SQL statements or programming is required.

Considering that the project specifications will not change significantly in the future, Intra.doc! was selected as the product that best met the needs of this particular project. Lindberg's

recommendation to use Intra.doc! has to do with the product's ease of implementation. An independent report on document management systems describes Intra.doc!'s implementation: "One of the key advantages of Intra.doc! is that it is an out-of-the-box solution that is extremely easy to use. Intra.doc! is easy to get up and running quickly, and the system uses simple HTML interfaces that users can access from any browser on any platform."<sup>1</sup>

## Project Implementation

Lindberg noted that the actual installation proved that it is possible to get Intra.doc! up and running in two to three days. After working through some problems with the NT Server OS the install of the Intra.doc! was accomplished in a matter of minutes. Working with the Intranet Solutions representative Kyle Hatlestad, Lindberg verified the choice of security plans and site taxonomy and moved ahead with building the post-install site. Before the end of the first day, the user accounts and security groups were set up and the directory structure was ready to be populated with documents. In fact the installation went so smoothly that Hatlestad and Lindberg were able to work on the first customization that first day.

One of the concerns has been the smooth integration of the Laboratory Overview Web site with the existing Director's Web site (<http://www.lanl.gov/worldview/news/director>).

Using the Intra.doc! Developers Kit (IDK), it was possible to customize the site, giving it the look and feel of the Director's site. The IDK was also used to simplify the search screen.

## Using the Web Site

The goal of this Web site is to provide Laboratory presenters with a point of easy access to presentation resources in the form of slides and multislides presentations. Both "Web viewable" (PDF) and "native" file formats (PowerPoint) are available.

## Steps to Access the Database

- Review the list of tools below and upgrade or install the following software: Web browser, Acrobat Reader or Exchange, and Microsoft PowerPoint.
- Start your favorite browser application and open [www.lanl.gov/laboverview](http://www.lanl.gov/laboverview).
- Find slide file or presentation files by "drilling down" through the directory structure or by using the Web site's search tool.
- Download/view a PDF version of an individual slide or multislides presentation with Acrobat plug-in enabled Web browser.
- If desired, download the native file for inclusion in your own PowerPoint presentation.
- To build a custom presentation using the Laboratory Overview slide files, download and use the PowerPoint template file: [laboverview.pot](#)

## Tools

To use this site you need a Web browser and the Adobe Acrobat plug-in and Reader. This site uses nested tables so problems may occur with older browsers. Recommended are Netscape Navigator (or Communicator) 4.X or Internet Explorer 4.X.

PDF files are made "Web viewable" through the use of Acrobat Reader (or Exchange) and a Web browser plug-in. Instructions on how to configure Adobe Acrobat Reader or Exchange plug-ins can be found at <http://eia.lanl.gov/pdfinstruct-main.htm>.

Netscape Navigator/Communicator and Adobe Acrobat Reader are available at no charge from <http://esd.lanl.gov>. Internet Explorer comes with both Macintosh and Windows operating systems.

To work with the PowerPoint native files you will need Windows version 97 or Macintosh version 98 or later.

## Finding Slides and Presentations

The system supports two approaches to locating documents. After accessing the site, you can browse (navigate) to a slide or presentation or use the search tool to find a document.

You can navigate to a slide or presentation by following the subject matter hyperlinks until you reach the link to the appropriate document. **Example:** To find one of the optional title slides by starting in the Laboratory Overview directory, you would click on the hyperlink to Title, org charts, budget, and statistics (see Fig. 1). This will bring up the Title, org charts, budget, and statistics directory page in your browser window. From this page follow the hyperlink to "Title Slides"; a Web page displaying links to all of the title slides will be loaded.

There are two ways to use the search tool to find a slide or presentation. Search for words within the document information (meta-data) or do a full-text search.



Fig. 1. Main menu of Laboratory Overview presentation resources.

Access the search tool by clicking on the magnifying glass in the tool bar at the upper-left of the Web page. The upper half of the page contains relevant document information fields.

**Example:** type "threat" in the Title field. Mouse-click either "Search" button—above or below the form. A page listing all documents with "threat" in their document information title will be listed. See Fig. 2 for an example of search results.

You may also perform a full search by scrolling down to the lower portion of the Search page and entering your search terms in the "Full-text" search window. **Example:** enter "threat" in the full-text field and click "Search." The results page will display all of the slide and/or presentation documents that contain "threat" in their text.

## Viewing the Files

To view the PDF version of the file, mouse-click the document icon or document title. If you have installed the Acrobat plug-in and configured your browser to use it, the PDF document will open in the browser's window. If you have not installed the Acrobat plug-in, the browser will either open the Acrobat Reader (or Exchange) application or write the PDF file to disk for you to open later with Reader or Exchange.

## Downloading the Native PowerPoint File

The link to the native file is below the document title on any of the pages that list documents or search results. A link to the native PowerPoint file is also found on the information page that can be accessed by clicking on the "I" icon to the right of the document title on the pages that list documents or search results.

Clicking the link to the native file will download a copy to your computer where—depending on your browser's configuration—PowerPoint will automatically be opened or the file saved on your hard drive for later access.

You may now include the slide in your presentation and modify it as necessary. For example, a title slide in PowerPoint may be modified to include the title of the presentation, presenter's name, and organization.

## Tips on Using PowerPoint to Create or Modify a Presentation

To create a presentation incorporating slides from the Laboratory Overview start by downloading a copy of the Laboratory Overview template: laboverview.pot. The link to this template is in the text at the top right side of pages that contain slide and presentation lists.

After downloading the template file and one or more of the presentation or individual slide files from the Laboratory Overview Web site, begin by opening the template file using PowerPoint 97 (4.X on the Mac) or later. In the PowerPoint menu bar go to "File – Save As" then select "a presentation" from the pull-down list and enter your own file name: <yourfilename.ppt>.



Fig. 2. Demonstration of search results showing thumbnail-size pictures of the slides.

## For More Information

If you have questions about the PowerPoint slides on the Web site, contact Ed Borrego at (505) 665-6257, or e-mail [eborrego@lanl.gov](mailto:eborrego@lanl.gov). For questions regarding the Web site or using *Intra.doc!* software, contact Lindberg at (505) 667-6884, or e-mail [cjl@lanl.gov](mailto:cjl@lanl.gov). If sufficient numbers of people are interested in training on using the Web site, Lindberg will deliver a course; contact him if you'd like to participate.

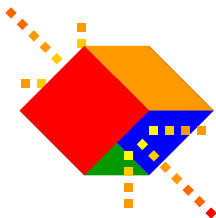
<sup>1</sup> Medina, Richard, and Kelley West, *Executive Brief: Intra.doc! Management System 3.6 from IntraNet Solutions, Inc.*, Doculabs, 1999, p. 12.

From the main menu bar go to "Insert," pull down to "Slides from Files . . ." Select the "Find Presentation" tab and then click the "Browse" button . . . Navigate through your machine's file structure until you find the files from the Laboratory Overview. (Netscape will save files to the desktop on the Mac or C:\Temp\ on a Windows machine by default.) Click "Open" once you have selected the file, then "Display." All of the slides in the presentation will be displayed in the "Slides Finder" window. Click on the slide you wish to import (hold down the shift key to select more than one file) then click "Insert." The slide will be added to your presentation.

The Slide Sorter view can be used to quickly move slides around in your presentation (note that on a Windows machine you will see the 4-slide icon in the lower left of your PowerPoint window). However, the slide sorter view should not be used to import slides from one presentation to another. The result will be a low-resolution image file that cannot be edited. (For information on beginning, intermediate, and advanced PowerPoint training, see this Web site: <http://198.59.97.19/UNMDats/UNMCourses.asp> .

## Desktop on Demand: A Desktop That Displays on Your Computer, but Actually Runs on the Enterprise Server

by Denise Sessions, BITS Managing  
Editor, CIC-6 Customer Service



Are you one of those people who move around the Lab and would like to use your own desktop no matter where you are? Do you need another computer to access Lab-wide or enterprise applications such as Time and Effort because most of your computing is done on a UNIX platform?

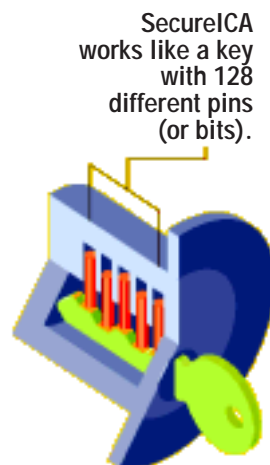
Business Information Systems, CIC-13, has created a great solution called Desktop on Demand for people like you who need remote access to either your own standardized desktop or enterprise systems. "Although this solution is not for every computer user at the Laboratory," says Michael Calhoun, CIC-13 Application Infrastructure Team leader, "some organizations can dramatically reduce the total cost of owning a computer by using Desktop on Demand."

To get the full story, BITS interviewed Calhoun and two other Desktop on Demand project team members, Jim Gore and Pat Hummer. Also to find out how Desktop on Demand can potentially

meet Laboratory customer needs, BITS polled usability testers who participated in a three-month beta or pilot program.

### Access from Windows, UNIX, and Mac

Desktop on Demand is an application available to PC, Mac, and UNIX users. The application resides on a remote server instead of on your desktop computer. The desktop that displays on your computer is actually a picture of your desktop. It runs on a centrally located cluster of terminal servers—the Enterprise Server—that users can access from any type of client (Windows, Macintosh, and UNIX). Your computer connects to the Enterprise Server via a well-encrypted Citrix Secure ICA® network protocol. Secure ICA software uses 128-bit encryption or key length to protect information being sent between servers running Citrix and clients (see Fig. 1).



**Fig. 1. Studies have proven that breaking an encrypted file of this strength would take millions of years and billions of dollars. Permission to use this graphic granted by Citrix Systems, Inc.**

## Enterprise Server Configuration

CIC-13's Enterprise Server runs on a cluster of machines (Compaq Proliant 6500 Quad processor Pentium II Xeon/400), each with 1.5 GB of memory. Each server is currently licensed for 115 concurrent users. The system software comprises the following packages:

- Windows NT 4.0, Terminal Server Edition
- Citrix Metaframe 1.8
- Systrack
- Compaq Insight Manager Agents
- Windows Scripting Host

While Desktop on Demand underwent beta testing, the Lab partnered with Intel Corporation to test the Enterprise Server with the new Intel Pentium III xeon computer chips. "The 500-MHz chips are faster than the current Pentium Pro computer chips on the server," says Calhoun. "The faster chips should handle more transactions from Lab employees." According to Calhoun, all new Enterprise Information Application (EIA) servers will use the Pentium III xeon chips. (For a related story about LANL testing Intel computer chips, see the July 12<sup>th</sup> issue of the *Daily Newsbulletin* at this URL: <http://www.lanl.gov/orgs/pa/News/071299.html>.)

## Long-Term Gains

The cost effectiveness stems from the long-term cost reductions associated with maintenance, hardware, and software. Desktop on Demand pilot user Bill Hargraves of ESH-5 said that his group wants to use Desktop on Demand to avoid having to buy new PCs and software upgrades every three years. He also likes the fact that no local backup is necessary because the Enterprise Server is backed up nightly.

Depending on a Laboratory organization's computing needs, only a monitor and terminal emulator are usually required to run Desktop on Demand. Given such simple requirements, Desktop on Demand dramatically decreases the need for desktop hardware support and individual machine maintenance.

With Desktop on Demand, organizations will be able to enjoy a new level of convenience by consolidating to one hardware platform. For example, it is not necessary to have another machine to do administrative functions if your main computing is done on a UNIX machine. You can simply log on to your Desktop on Demand by clicking on the link on the EIA Desktop Web page. "It couldn't be easier," according to CIC-2's Jack Harris, who is working with Sun and Silicon Graphics users who are using Desktop on Demand in X-Division. Among the advantages these UNIX computer users are enjoying include using their current printers in UNIX and storing and retrieving files on UNIX file servers. They are using applications such as Eudora and MS Office, and enterprise systems such as Travel and Time & Effort while saving space in an environment where desktop space is at a premium. In fact, these users save time by not having to learn to use a new machine for secondary computing needs.

## Front-End Costs

Costs and recharge fees include

- a monthly access charge of \$29.95
- a monthly storage fee of \$14.15 per GB
- NT client access license, a one-time \$81.00 fee (for non-NT computers) (Note: NT versions of commercial software are available on the Lab's electronic software distribution site, ESD.)

Users will also need Netscape version 4.0 or higher, or Internet Explorer 4.0 or higher. The Citrix ICA Client, a special communications software, is free.

## Best of All Possible Worlds

Because Desktop on Demand provides easy access to enterprise systems and commonly purchased software, users don't have to wait or use disk space to download entire enterprise applications, such as Data Warehouse. Also, the project promotes Information Architecture (IA) software standards by offering a standardized desktop and has established a common set of desktop software tools for use throughout the Laboratory. All software applications you would typically have on a stand-alone computer are equally available on Desktop on Demand, such as Terminal Emulation, MS Word, and Excel. Even though the desktops are set up with NT versions of IA-standard software, the Desktop on Demand project team can customize desktops to accommodate the unique needs of specific organizations.

Also, in keeping with the IA software model, organizations such as HR-5's Staffing Alternatives Team can improve productivity and reduce training costs by training employees to work with standardized software. Delilah Garcia, a Staffing Alternatives Team member, says that Desktop on Demand was well received by their employees whose job assignments and locations change frequently as they respond to fluctuating workload demands around the Laboratory.

Desktop on Demand offers users important security benefits. Desktop on Demand uses industry-standard encryption. All information transferred between the user computer and the server don't actually contain data in a useable form. Your desktop is stored on the Enterprise Server, which is backed up nightly. You can also store your data files on the server. The files on the server are substantially more stable than on your own hard drive because the files are on a protected fault tolerant disk system.

## Accessing Desktop on Demand

For those who could benefit from this new EIA application, Desktop on Demand is available through the EIA Web site at <http://enterprise.lanl.gov>. If you have problems or questions, call Enterprise Information Applications Consulting at 5-4444, option 2.

## Grassroots Software Management: Simple Things to Do

by Don Willerton, Group Leader,  
CIC-6, Customer Service

### Software Project Management Series

In the previous two articles, we've talked about managing software projects. The first article focused on how traditional project management techniques fall short of addressing the inherent uncertainty in software development. The second article looked at three examples of increasingly nontraditional methodologies for addressing projects with high uncertainty. This last article boils down some of current themes and suggests some practices to increase the success rate for software projects.

Gleaning the advice from several different authors is fun. Definite patterns of beliefs show up and, more often than not, most good advice tends to be simple and straightforward. To be more successful in building software, do we have to adopt huge methodologies that require lots of time, effort, and money? Do we have to become converted to a "new way of life" and become zealots of a new order? Are the answers in forming massive, complex, matrixed teams that satisfy every conceivable process and requirement?

Well, if you are going for above the 90th percentile for guaranteed project success, maybe so.

But tremendous gains can be had for a lot less time, effort, and money and improve your probability of success considerably. That level of increase may be sufficient for your team or organization. The following is a simplified list of good practices or themes identified by the Software Engineering Institute, Steve McConnell, the Airlie Council, the Adaptive Project Management method, Tom DeMarco, Rob Thomsett, Alan Davis, and other authors. Most of these practices are just common sense. I've grouped them to make it easier to remember.

### People Are the Most Important Resource

#### Hiring

The most significant factor in increasing project productivity is to have very good people.

#### Antifragmentation

Project members who can focus for significant periods of time and who are not switching back and forth to other projects, will be far more effective contributors. (At least four authors identified fragmentation as the number-one enemy of projects.)

#### Wisdom and Courage as well as Skill

Uncertainty in projects requires people who can see dilemmas, contradictions, inconsistencies, and conflict, and yet have the wisdom and courage to work comfortably in the environment. Uncertainty is no place for wimps.

#### Dedicated Small, Permanent, Senior Core

Several authors emphasized the benefits of having good, proven, experienced people on a project for the entire duration, especially from the viewpoint of corporately learning to do projects better.

#### Leadership

The project manager is a prime, if not the prime, focus of what the project is, what it does, and what it produces. They must be skilled *managers*, not just technologists.

#### Match Project to Interests

People do better at doing what they're interested in. Learn to reframe project activities to match people's personal interest.

#### Teams

Especially in an environment of high uncertainty, high risk, and high complexity, isolated individuals can not do projects. With an existing team, great care must be used in managing the team dynamics.

### Learning Is Not a By-Product, It Is an Objective

#### Learning environment

Everyone needs to have the mindset that learning how to do the project is a significant goal in a project. No project lends itself to a formula anymore.

#### Post Mortems

This is one of the most important practices. Post mortems must be periodic, formal, complete, and open, and *must* include the customers.

#### Nothing Ad Hoc

Increased cycle time doesn't mean not being disciplined. Disciplined actions, recognized and followed by the whole team, increase learning and adapting.

## Make and Find Your Mistakes Early

### Reviews

You must have design reviews, code reviews, test reviews, documentation reviews, etc. One design error found in maintenance phase will cost 50 to 200 times the cost as when caught in the design phase.

### Timeboxing

This is a mechanism that encourages periodic releases and post mortems, high visibility for the whole project community, team learning, and customer involvement.

### Rework

"Rework" is work that's done over again and includes fixing bugs, redesigning interfaces, changing for new functionality, recoding for requirements change, etcetera. Rework may be the most costly of all tasks that is almost totally ignored by current common practices. Instead, it must be recognized, estimated, and included in planning.

## Make Little Mistakes

### Staged Delivery

Get something into the users' hands as soon as possible, even if it's a partial product. Get used to involving them, listening to them, and doing what they say. This may cause requirements to change and work to be done over, but the product will be significantly better in the long run.

### Minimalism

Create products, subroutines, functions, or elements that represent the minimal implementation of their function, no more, no less. Minimize the dependencies between modules or functions.

## At Least Don't Make the Same Mistakes

### Software Reuse

The most efficient way to produce software is to not have to produce it at all. This has the greatest potential for automatically creating products faster, cheaper, and of higher quality.

## Customer Are Your New Best Friends

### Project Identity

Customers who are involved have higher identity, more sympathy, make better contributions, add more excitement, and will be better resources when they are treated as collaborators.

### Continual Involvement in Reviews

Use customers, stakeholders, managers, and others to increase the corporate wisdom of the project, as well as sharing the adapting that takes place. Everyone gets to confront the tough issues.

Some of these practices are well known and have been around a long time. Some are newer and more innovative. They are all trying to improve the development of software in an increasing volatile, uncertain, and competitive environment. If our desires are to increase the fun, satisfaction, and effectiveness of projects, as well as continually improving our success at doing projects, these are worthwhile practices to learn.



## News About Microsoft Networking at Los Alamos: The LANL Master Accounts Domain Is Up and Running!

by Cheryl Host, NT-MAD Project Leader, CIC-5, Network Engineering

The LANL Master Accounts Domain (MAD)—a move to centralize Windows NT user accounts—is running in pilot mode with over 950 user accounts. The Business Information Systems group, CIC-13, is using the LANL domain for authentication for Desktop-on-Demand and other Laboratory-wide applications. Soon, the Electronic Software Distribution (ESD) program will use the LANL domain for its accounts. The NT-MAD project is piloting a program with JCNM to migrate accounts in existing domains to the LANL domain. NMT Division is establishing its new NT network to use the LANL domain, and ESH Division plans to migrate to the LANL domain in September.

### Advantages

The advantages of this consolidation of networks are

- a single logon gives users easy access to many resources, for example enterprise or Lab-wide applications like Travel and Time & Effort, and group or project file shares, printers, and applications;
- resource domain administrators can give users access to additional resources simply by adding the user to the resource access control list; no need to set up a specific user account or additional trusts, or worry about name resolution issues;
- groups/projects/divisions can share data and applications more easily;
- resource domain administrators can more easily accommodate users who move between offices;
- security audits are simplified because resource domain administrators perform account maintenance and security at the LANL domain level rather than at the local domain; and

- as improved security measures are applied to the LANL domain, overall NT security at the Lab is enhanced.

### Quick Access to the LANL Domain

Individuals can create LANL domain accounts using the Web-based Integrated Computing Network (ICN) registry at <http://www.nic.lanl.gov/register>. After the account is created, the new users should be able to see "LANL" as a domain option on their login screen (see Figure 1). If they don't see the domain option, they are instructed to contact their system administrator.

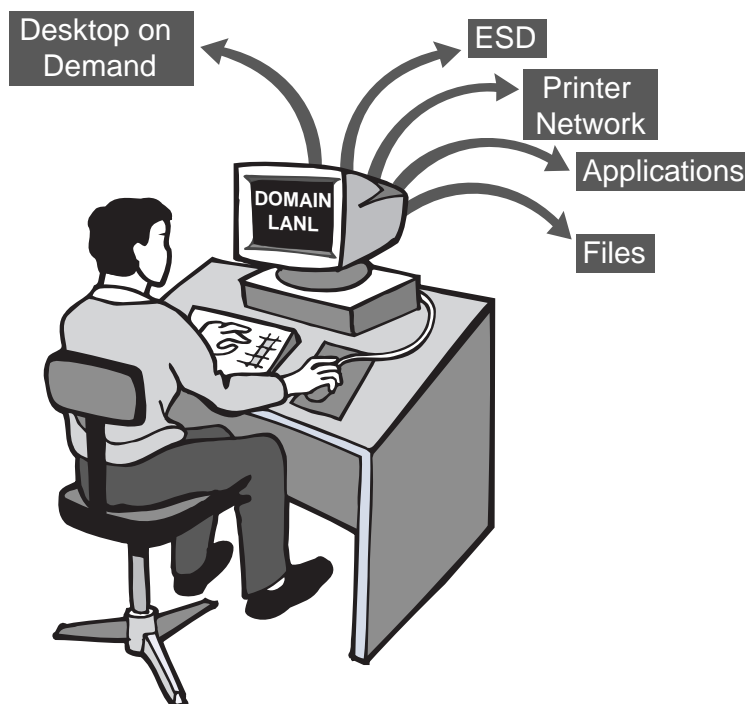
### Additional Information

You can get additional background information on the LANL MAD project at <http://www.lanl.gov/projects/ia-lanl/area/infra/nt/>.

### Notes for NT System Administrators

#### Setting up a Trust between Local NT Domain and the LANL Domain

Contact Cheryl Host, [chost@lanl.gov](mailto:chost@lanl.gov), or Bob Stewart, [res@lanl.gov](mailto:res@lanl.gov), to set up a trust between your local NT domain and the LANL domain. This procedure requires a short telephone call or e-mail exchange. Setting up the trust takes only a few minutes and all we need is your domain name. We then add your domain as a trusting domain and enter a password for set-up purposes. You then go into "User Manager for Domains / Policies / Trust Relationships..." under Administrator Tools and enter LANL as a trusted domain and the password we give you. Within about a minute, the trust is established.



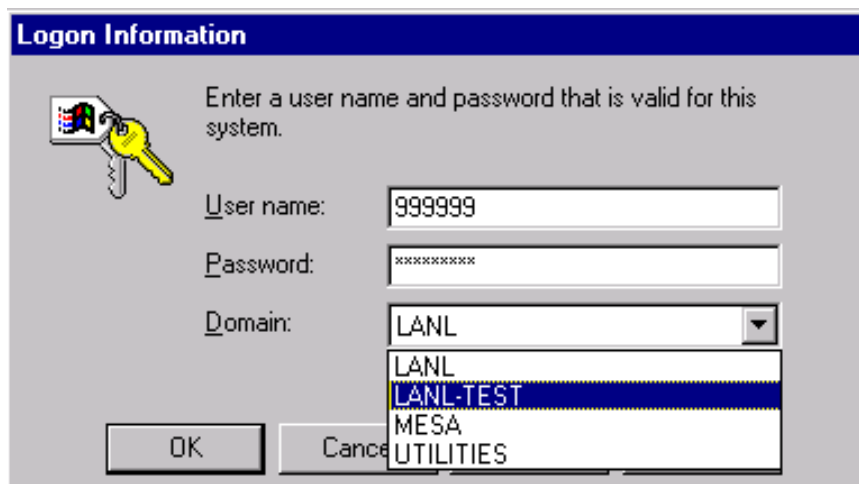


Fig. 1. An example of the LANL MAD login screen.

This process only has to be done once per domain. Then, any other member of that domain who sets up a LANL domain account will be able to see and choose the LANL domain as an option.

### Setting Up a New Domain

If you are establishing a new NT domain, it is much easier to bring it up under the LANL MAD model because it eliminates the need to migrate Access Control Lists (ACLs), profiles, and other account information. Please contact Cheryl Host or Bob Stewart for information on how to begin.

### Migrating a Large Local Domain

If you have an established domain to migrate into the LANL domain, we have tools to ease the migration of many accounts. The first step for migration is to ensure user name compatibility. The LANL domain uses Z-number for user names. User names that include Z-number are more easily migrated.

### Centralized WINS Servers Available

LANL has WINS servers available for the Laboratory population. Centralized WINS services provide accessibility (that is, browsing across subnets) to Microsoft systems throughout the Laboratory. There are over 3000 systems registered with these WINS servers.

The WINS server addresses are

- 128.165.5.5 (primary),
- 128.165.5.6 (secondary), and
- 128.165.155.165 (Administrative partition).

The "senate" machine, 128.165.129.1, is being phased out. Service will be shut off on September 30, 1999. Many people had been using this machine for centralized WINS services. Please convert your WINS settings before this date.

### Avoiding Naming Conflicts

Microsoft systems currently have a naming scheme that is independent of the Internet DNS naming convention. This inconsistency means that systems can have one Microsoft networking name and a different DNS name. CIC-5 strongly encourages administrators to choose the same name for a machine's networking and DNS name, and register that name with the LANL Hostmaster. Please see <http://protected/nst/hstmstr.html> for Hostmaster information.

Future versions of Microsoft operating systems will use DNS names, so it is to your advantage to ensure they are in synch now.

Microsoft domain names must also be unique across the entire network to ensure communication among Laboratory systems. If you are a Microsoft resource domain administrator, please register your domain name by sending e-mail to Cheryl Host, [chost@lanl.gov](mailto:chost@lanl.gov).

## Notes on Setting Up and Configuring an NT Domain

by Andy Ryan, Microsoft Certified Systems Engineer, EES-13, Integrated Geosciences Group

### Correction

In my last article<sup>1</sup> about minimizing NetBIOS traffic on the network, some of you may have noticed that I said when the Windows workstation node type is hybrid (0x8), the ARP cache and WINS server were queried for NetBIOS name resolution. What I wanted to say is that the **NetBIOS Name Cache** and the WINS server were queried. The NetBIOS Name Cache is the ARP cache's counterpart in the NetBIOS environment.

### Overview

When you set up an NT workstation, it becomes a member of a workgroup or a domain. For networks with less than 20 workstations, a workgroup environment is recommended. This means that each user is responsible for administration of their workstation in a peer-to-peer environment. As you probably know, this scenario presents account/password synchronization issues. And although you may designate a workstation as a "server", the user accounts remain distributed among the workstations.

To ease and centralize administration, enter the NT Domain. When a workstation joins a domain, it has access to a central accounts database that resides on specialized servers called Domain Controllers (DCs). The benefits of the domain environment are many. Some of the benefits include centralized account and resource administration, decreased network traffic, added security, and

login script capability. Another advantage is the ability to use remote administrative tools that do not work well in a workgroup environment like User Manager for Domains, Server Manager, Event Viewer, Command Scheduler, and Policy Editor.

### Architecture

An example NT domain architecture is presented in Figure 1. The pictured NT domain architecture integrates well into a subnetted TCP/IP environment. The example configuration provides robustness, scalability, maintainability, and ease of administration. Basically, the simplest NT domain consists of a Primary Domain Controller (PDC) and a Backup Domain Controller (BDC) on a single subnet. There can be an infinite number (hypothetically) of BDCs per domain, but only one PDC. As the NT domain expands to additional subnets, a BDC (with WINS) should be located on each additional subnet to keep Microsoft networking traffic through the routers to a minimum and to optimize logon and lookup speed.

### Requirements

All that is needed to implement the basic NT domain are two PCs, two copies of NT Server, and a client access license for each user account. One computer serves as the PDC and the second as the BDC. Also, you should allow one additional BDC for each additional subnet. Using this proposed domain architecture, the number of domain controllers can be found by using the formula,  $2 + (n-1)$ , where  $n$  is the number of subnets.

In a recent configuration with three subnets and about 500 accounts, I used Dell PCs with a Celeron 400 CPU, configured with 128 MB of RAM, a 4-GB IDE disk, and 100baseT. According to the white paper, "Domain Sizing and Capacity Planning for Windows NT Server 4.0," which can be found on the Technet CD<sup>2</sup>, my hardware is overkill for 500 users. But the Celerons were the most economical machines available. For 40,000 users or more, the white paper recommends SMP on the PDC, although you may want to look at SMP sooner than that. The

### Example NT Domain

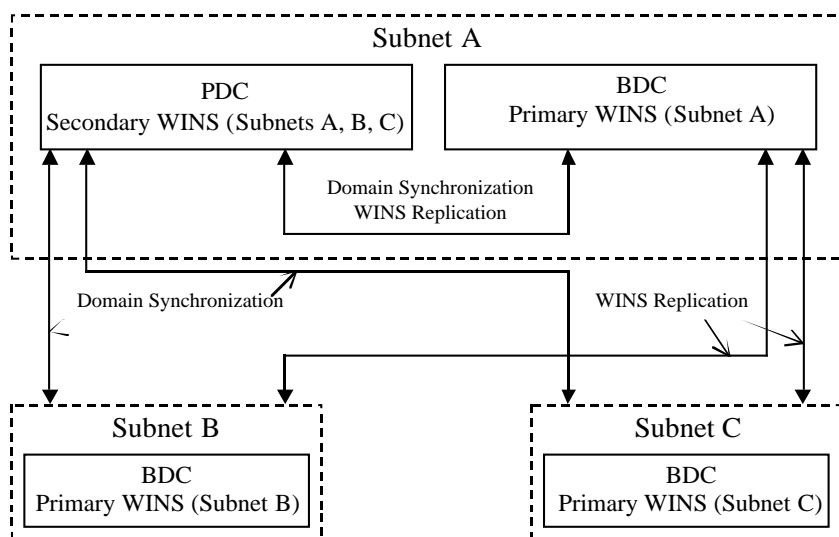


Fig. 1. An example NT domain architecture.

reason for SMP is that having more than one CPU on the PDC allows external requests to be processed without interrupting the domain synchronization process.

Most likely, redundant hardware will not be needed on the DCs because the restoration process is usually a reinstall of NT Server. But you should backup the registry of the PDC on a nightly basis. You can use REGBACK.EXE and WINAT.EXE from the NT Resource Kit to schedule backups of the registry. Having a current copy of the PDC's registry provides a means of domain restoration in the event that all DCs fail concurrently or a domain SID becomes corrupt or damaged. Also, if you are running services that use database files, like WINS and DHCP, you should backup the databases regularly.

## Installation

The installation of the PDC must be performed first, because the domain is created by bringing up the PDC. Choose a name for your domain that you can live with because changing the domain name is to be avoided at all costs. Install NT Server on each computer as you normally would, using the following notes as a general guide.

Delete all disk partitions and cut a  $\geq 2$  GB NTFS partition as the C: drive and a  $\geq 500$  MB partition for the paging file. The minimum recommended paging file size is RAM + 12 MB. Make sure the paging file partition is on the same physical disk as the NT system. Further, it is best that the paging file not be allowed to grow during the operation of the system because it adds time to the paging processes and causes undue fragmentation. To keep the paging file from growing, just set the initial and maximum sizes the same.

During the first installation, create the NT Domain by choosing **Primary Domain Controller** at the appropriate setup window. You'll want to hardcode the TCP/IP configuration for reliability. Also, enable LMHOSTS lookup and maintain an LMHOSTS file that contains all DCs and WINS servers. A sample LMHOSTS might read:

```
128.165.180.25 MYPDC
#DOM:MYDOM #PRE
```

```
128.165.180.26 MYBDC
#DOM:MYDOM #PRE
```

```
128.165.180.27 MYWINS #PRE
```

```
128.165.100.200 SUBBDC
#DOM:MYDOM #PRE
```

```
128.165.100.201 SUBWINS #PRE
```

Install the WINS, DHCP, and DNS services during setup as appropriate. If you use any of these services, it makes good sense to run them on the DCs. The DCs then provide a logical grouping of core services and keep traffic like name registration and resolution, logon validation, domain synchronization, DNS lookup, and browsing clear of other servers. A typical DC might run:

DHCP

WINS

Browsing Services

DNS  
Domain Directory Services

Licensing Management Services

After bringing up the BDC(s), use Server Manager to synchronize the domain. This forces the PDC to send a current copy of the domain database to the BDCs. When the PDC and BDCs are synchronizing correctly, which can be determined from the system log or by typing the command

```
nltest /bdc_query:DomainName,
```

You are ready to migrate accounts from local machines to the domain account database and join workstations to the domain. My next article will talk about moving your NT workstations from a workgroup to a domain environment. One would think that this would be a simple, straightforward task, but there are some traps and pitfalls. Unfortunately, Microsoft has offered a NetWare to NT migration utility, but not an NT workgroup to domain migration utility.

## Optimization

By default, DCs have their server service configured for **Maximize Throughput for File Sharing**. While this is correct for a file or print server, it does not provide the best performance for a DC that needs to validate logon requests. Instead, configure the server service of all DCs for **Maximize Throughput for Network Applications**. By properly configuring this option, the number of simultaneous logons can be increased from about 6–7 per second to around 20.

Further, you may tune the WINS, DHCP, NetLogon, and other services according to your needs and environment via the management programs and registry, although the default settings work for most applications. Article Q142692, "Minimizing WAN Traffic", and Document NET405EF.DOC, "Network Traffic Analysis and Optimization", both provide excellent information about tuning network services. These documents can be found on the Technet CD.

## Summary

When the number of Windows workstations on your network approaches 20, you should consider implementing an NT domain. The main benefit of an NT domain is ease and centralization of administration. Planning and following rules are key in setting up an NT domain.

First, design your architecture by determining how many DCs you will need, where they will be located, and what services they will provide. Next, create the NT domain by bringing up the PDC. Then bring up the BDC(s) and verify domain synchronization. Last, tune the networking service(s). You should now have a functional NT domain that is ready for account and workstation migration.

<sup>1</sup>See BITS February-March 1999, p. 8, at <http://www.lanl.gov/cic/publications.html>.

<sup>2</sup> For information about Technet see <http://www.microsoft.com/technet/subscription/how.htm>.



## Web Content Architecture: Taming the Tangle of Protocols

by Tad Lane, Information  
Architecture Standards Editor, CIC-1  
Communication Arts and Services

*Note:* This article is adapted from a presentation given at the IntraLab99 Conference for Web developers. Slides from that conference are available online at <http://www.lanl.gov/www-team/seminars/intralab99.html>.

Early in Marquez's *One Hundred Years of Solitude*, there's a scene where Jose Arcadio Buendia is working with his sons in their alchemy laboratory. The boys "grew enthusiastic over the flying carpet that went swiftly by the laboratory at window level carrying the gypsy who was driving it and several children from the village." Jose Arcadio Buendia ignores the flying carpet, saying "Let them dream.... We'll do better...with more scientific resources than a miserable bedspread."

Alchemy vs. flying carpets. It does seem applicable to the Web. On one hand, we've got the Philosopher's Stone of open standards; on the other, there's the flying carpet of proprietary vendor extensions.

The real alchemist's trick here is that even as open standards are appearing to get more complicated, they are actually making things simpler. HyperText Markup Language (HTML) has grown into eXtensible Markup Language (XML), eXtensible HTML (XHTML), Document Object Model

(DOM), Cascading Style Sheets (CSS), with scripting, multimedia, database access, and more tossed in for good measure. But the result is simpler because it allows each specification to focus on a specific area, instead of trying to make a single language accommodate it all.

In this article, I'll lay out my best guess at how the various specifications interrelate. It's a high-level view, not looking into the details of any particular area. Hopefully, though, the "big picture" can help each of the parts make more sense.

Vendor implementation of each of the parts is beyond the scope of this article. Some parts, such as XHTML, are nearly fully supported by leading browsers. Others, such as CSS are supported inconsistently at best. What's worse, anything we say today about the state of implementation will be out of date as soon as the next browser upgrade comes out. To focus on the theory behind it all, this article will simply avoid the question of implementation.

### Overview

The overall technical architecture of the Web can be separated into three broad areas:

- Rendering Architecture: The browsers and other "user agents" that interpret Web content.
- Content Architecture: HTML and other languages in which Web content is written.
- Transport Architecture: The underlying protocols that control the electronic connections.

The rendering level includes a lot more than just traditional browsers. We already have a variety of indexing robots, which read and interpret Web content automatically. There are plug-ins for special needs, voice browsers, refreshable Braille pads, wireless palm units, and more, all needing to access the same Web content and providing the fundamental reason to design Web content for accessibility. It's also due to the variety of these tools that "user agent" is generally the preferred term to describe them, regardless of whether we're referring to a traditional browser, a voice browser, an indexing robot, or whatever.

The transport level is what controls the actual transmission of electronic signals between the server and the client. This includes the basic "Transmission Control Protocol/Internet Protocol" (TCP/IP) suite, along with Web-specific protocols such as HTTP and Secure Sockets Layer (SSL) (which provides encryption).

In between the rendering and transport levels sits the Web content which is transmitted across the transport level and interpreted at the rendering level. Although the levels are distinct, they have interrelated issues. For example, the content formats should be as lightweight as possible to minimize the volume of traffic at the transport level, and simplify the interpretation of the content at the rendering level. Not only does a bloated HTML increase the volume of traffic (by requiring, for example, repetition of formatting instructions), but it also leads to large, complex browsers (because the

software needs to be programmed to understand every attribute for every tag, etc.). An improved content architecture should correct the bloating of HTML.

## The Evolution of Content Languages

Traditionally, all markup of Web pages has been done with one language, HTML. As illustrated in Fig. 1, HTML itself is written in Standard Generalized Markup Language (SGML), which can be thought of as a language for writing other markup languages in.

User Agent
HTML
SGML
Transport Protocols

**Fig. 1: Traditional Markup Language Model**

On the face of it, this approach seems simple enough, but it has two fundamental problems. First, as its name implies, SGML is a generalized language that has to accommodate a whole range of different needs that have nothing to do with network communications. For example, SGML can be used to instruct a printer about the formatting of a hardcopy page. As a result, SGML is a large, complex language that permits certain inconsistencies, such as sometimes requiring tags to be closed, sometimes requiring tags to not be closed, and sometimes leaving the closing optional. HTML 4.0, written in SGML, is full of such inconsistencies, and user agents need to be programmed to understand them all.

The second problem is the bloating of HTML mentioned above. As we try to do more things with the Web, if we try to support everything with just one markup language, then that language becomes increasingly large and complicated. For example, to control text alignment, add a <center> tag; to control text color, add a <font> tag; to accommodate scripting, add a <script> tag; to include client-side image maps, add a <map> tag; to bring in Java, add an <applet> tag; etc. As a result the markup language bloats, network traffic goes up, and user agents get larger and more complicated.

As illustrated in Fig. 2, the emerging markup language model appears at first to be more complex. It is actually simpler, however, as explained below.

User Agent			
XHTML	Math ML	SMIL	etc.
XML			
SGML			
Transport Protocols			

**Fig. 2: Emerging Markup Language Model**

The first thing to note is that XML has been placed between SGML and the other languages. XML is a specialized markup language that has been optimized for network communications. Like SGML, it is a language for writing other markup languages in, but unlike SGML, it is simple and internally consistent (e.g., all tags must close). Note that it is also possible to write pages directly in XML, but that its most significant use in terms of the overall model comes from writing other markup languages.

Above XML are a series of markup languages written in XML, each focusing on a specific need:

- XHTML is the next generation, XML-compliant version of HTML. It is a general purpose markup language, written to do the same things that HTML currently does (except more consistently).
- MathML is the Mathematical Markup Language, which provides complex mathematical notation.
- SMIL is the Synchronized Multimedia Integration Language, which enables us to pull together text, sound, motion pictures, and images into a TV-like presentation.

Other XML markup languages address other specific needs. Examples include Chemical Markup Language (CML), Scalable Vector Graphics (SCG), Bioinformatic Sequence Markup Language (BSML), Genealogical Data in XML (GedML), and more.

What's simple about this approach is that user communities can develop the language and tools for their own applications without needing to worry about other applications. If I want to use MathML, I don't need to worry about GedML. Generally, specialized software is needed to fully understand a particular language, but the simpler languages lead to simpler software, which is quicker to load and run. Instead of one language trying to do everything and becoming bloated in the process, we have targeted languages that do fewer things but do them better.

## Content Components

The markup language itself is only a portion of the overall Web content architecture. There are also images, scripts, multimedia, etc. As shown in Fig. 3, these are frequently in external files that are pulled into a document through the markup language (e.g., an `<img>` tag).

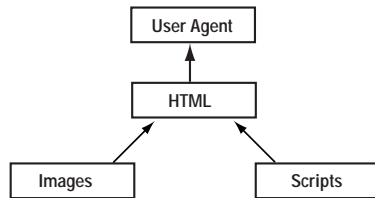


Fig. 3: Traditional Content Components

Again, this looks simple enough at first glance, and it works well enough as long as pages are simple enough. Web documents are no longer just text plus images plus simple scripts. Instead, they're evolving into a mixture of text, images, scripts, multimedia, updateable components, and more. Toss in advanced formatting, document definitions, and digital signatures, and the traditional approach can't sustain the weight.

As with the markup languages themselves, the emerging component model appears at first to be more complicated, as illustrated in Fig. 4. Again, though, it is actually simpler.

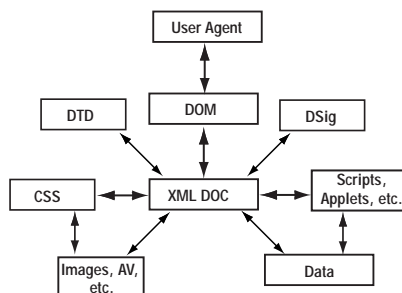


Fig. 4: Emerging Content Components

As with the markup languages, there are more components in the emerging model. And just like the various XML languages focusing on a targeted application, each of the components in the new model focuses on a particular task, enabling to perform that task in a better, more streamlined way.

As with the HTML document in the traditional model, the core XML document sits in the middle of the overall component model. It pulls in the various other components via markup such as `<img>`, `<object>`, or `<link>`, and it serves as a kind of assembly area for constructing the full document. Arrows in this model now go both ways, to reflect the increased interactivity of the emerging model. Among the various components that the XML document can pull together are the following:

- The Document Object Model (DOM) defines the interface between the user agent and the Web content. It defines how, for example, to identify the third cell in the fourth row of the second table, so that programs and scripts can dynamically update it.
- The Document Type Definition (DTD) defines the XML language that the document is written in. It explains to the user agent what the content of the document is (e.g., `<h1>` is a first level heading, `<li>` is a list item). A DTD has always existed for each prior version of HTML, but its definitions have previously been hardcoded into the user agent. Under the emerging model, the DTD needs to be separate so that it can be changed from one XML language to the next.
- The Cascading Style Sheets (CSS) define how the document content should be rendered. For a graphical browser, for example, `<h1>` might be rendered in centered, bold, large font. A number of HTML tags and attributes that have previously done this will hopefully soon be superseded by style sheets, including `<center>`, `<font>`, and similar rendering tags. The user agent can substitute a different style sheet if the user prefers, which would enable, for example, a voice browser to speak "level one heading" for `<h1>`.
- Images, video, sound, etc., will generally remain in external files, just as they have traditionally. Some of these files might themselves be written in an XML language, however, such as Scalable Vector Graphics (SVG).
- Scripts, applets, and similar programs can be executed on either the server side or the client side, and the data that they access can be on the server or client as well. What the model is meant to show is that data will be more easily accessed and updated, leading to a more dynamic and interactive environment. Instead of needing to fill out a full Web form and then press submit, for example, the new model enables data to be updated as it is entered. This can lead the way to applications such as Web "whiteboards," where a group of coworkers can work on the same problem at the same time and see what the others are doing as they do it.
- A Digital Signature (DSig) is a cryptographic hash that can validate that the document has not changed since person "A" electronically "signed" it and asserted it was valid. For an official document, for example, the DSig can show that the document has not been hacked. The signature needs to be external to the document, since adding the signature into the document would change the document, thereby making the signature invalid.

## Acronyms and Such

*CSS—Cascading Style Sheets (W3C<sup>1</sup>)*

*DOM—Document Object Model (W3C)*

*DSig—Digital Signature (W3C)*

*DTD—Document Type Definition*

*HTML—HyperText Markup Language (W3C)*

*HTTP—HyperText Transport Protocol (IETF & W3C)*

*IETF—Internet Engineering Task Force*

*IP—Internet Protocol (IETF)*

*MathML—Mathematical Markup Language (W3C)*

*SGML—Standard Generalized Markup Language*

*SMIL—Synchronized Multimedia Integration Language (W3C)*

*SSL—Secure Sockets Layer*

*SVG—Scalable Vector Graphics*

*TCP—Transmission Control Protocol (IETF)*

*W3C—World Wide Web Consortium*

*XML—eXtensible Markup Language (W3C)*

*XHTML—eXtensible HTML (W3C)*

<sup>1</sup>Standard issued by W3C and/or IETF.

This is, as usual, not a complete picture of all the components that are evolving. There is also, for example, Extensible Stylesheet Language (XSL), which offers a different type of style control for nonlinear, dynamic documents (such as automatically adding a table of contents at the time a document is requested). Another World Wide Web Consortium (W3C) working group is looking into XML signatures for nonstatic documents (enabling, for example, a portion of a document to be signed while other parts are not). With similar efforts, additional content components will continue to emerge. Hopefully, though, the simplified model presented above can help to make the interrelationship of the various components easier to understand.

## For Further Information

The following BITS articles have addressed related subjects:

- Extensible HTML: Preparing for the Next Phase of Web Markup (August 1999)
- Extending Web Documents: Getting Ready for XML (March 1998)

For additional information about Information Architecture (IA) Web activities, please see our Web team page at <http://www.lanl.gov/projects/ia-lanl/area/web/>. For additional information about the IA Project in general, please see our project home page at <http://www.lanl.gov/projects/ia/>.

## Using Tecolote Components to Extend Object-Oriented Programming

*by Mark Zander, Technical Staff Member, CIC-12, Scientific Software Engineering Group*

Tecolote was written to support Blanca physics applications, as described in the K. Holian article in the last BITS issue.<sup>1</sup> (The Blanca Project supports the goals of the Accelerated Computing Initiative at Los Alamos National Laboratory.) Tecolote, a single software support layer for Blanca applications, has its own unique component architecture.

Tecolote's component architecture is designed to extend traditional object-oriented (OO) programming in C++. We had found that although OO allowed us to write functions that are independent of specific objects on which they operate, OO did not permit us to write complete programs that are independent of specific objects. Therefore, we developed a component architecture that allowed users to defer explicit specification of objects until run time so that programs can be independent of the objects they use.

### What Is Object-Oriented Programming?

Object-oriented programming is a method of creating programs organized as collections of cooperative objects. An object is a grouping of data with its associated functions. When separate objects share the same layout or description they are instances of the

same class. Objects with different descriptions may exhibit the same behavior using a technique called inheritance. This allows for functions to be written that are independent of the type of objects they manipulate.

A class is a description of objects. It consists of data declarations and functions that manipulate that data. The data associated with a class is referred to as the member data of that class and determines the layout of objects of that class. The functions associated with a class are referred to as methods of that class and determine the operations to be performed on the class member data.

The C++ class circle listed below is defined in four steps. First the class header lists the class name and the base class from which it inherits. Second, this is followed by the data members of the class. Third, a special constructor method is defined that sets all of the data members to their initial values. Fourth, methods are defined that manipulate the member data.

```
// 1. Make a class circle that inherits from base-class shape
```

```
class circle: public shape {
```

```
// 2. Define data members of class circle
```

```
float    x_center, y_center;
```

```
float    radius;
```

```
public:
```

```
// 3. Define a special constructor function to initialize data
```

```
circle(float x, float y, float r)
```

```
: x_center(x), y_center(y), radius(r) { }
```

```
// 4. Define methods (functions) of class circle
```

```
float area() { return 3.14 * radius * radius; }
```

```
bool is_inside(float x_pt, float y_pt) {
```

```
return radius * radius <
```

```
pow(x_pt - x_center, 2 ) + pow(y_pt - y_center 2 );
```

```
}
```

```
};
```

An object is an instance of a class. For each class in our system, we may have several objects representing that class. Each of these objects has the same types of data and the same functions, but different values may be present.

In the following C++ code, we demonstrate how to create two objects from class circle. The first object, c1, is a circle centered at (0.0, 0.0) with a radius of 5. The second object, c2, is a circle centered at (1.0, 1.0) with a radius of 2. After creating the two objects we print out their respective areas.

```
circle    c1(0.0, 0.0, 5.0);
```

```
circle    c2(1.0, 1.0, 2.0);
```

```
cout << "c1 area = " << c1.area() << endl;
```

```
cout << "c2 area = " << c2.area() << endl;
```

Inheritance relationships allow objects represented by different classes to have the same behavior. A class that defines general properties attributed to several more specific types is called a base class. A class that implements

these special properties for a specific case is called a derived class. Circle is a derived class of shape and shape is the base class for circle.

In the C++ class shape below, we declare two pure virtual methods that will need to be implemented in classes derived from class shape. Shape supplies an interface definition for all of its derived classes.

```
class shape {  
    virtual float area() = 0;  
    virtual bool is_inside(float x_pt, float y_pt) = 0;  
};  
  
We may define a second C++ class, box, which is also a shape:  
  
class box : public shape {  
    float    x_start, y_start;  
    float    x_size, y_size;  
public:  
    box(float x, float y, float x_s, float y_s)  
    : x_center(x), y_center(y), x_size(x_s), y_size(y_s) { }  
    float area() { return x_size * y_size; }  
    bool is_inside(float x_pt, float y_pt) {  
        return (x_start < x_pt && x_pt < x_start + x_size) &&  
        (y_start < y_pt && y_pt < y_start + y_size);  
    }  
};
```

We may now write code that is independent of the type of shapes. For example if we wish to write a function that tests if a point is inside two shapes we may write:

```
bool intersection_point(float x_pt, float y_pt, shape* s1, shape* s2) {  
    return s1.is_inside(x_pt, y_pt) && s2.is_inside(x_pt, y_pt);  
}
```

In order to write code that uses "intersection\_point" we must be explicit about which shapes are being used:

```
main() {  
    circle    c(1.0, 1.0, 0.5);  
    box       b((0.0, 0.0, 1.0, 1.0);  
    bool      test;  
    test = intersection_point(1.0, 1.0, &c, &b);  
    cout << "test = " << test << endl;  
}
```

### What Is a Component Architecture?

In the above example we wrote a **function** that is independent of the specific objects it uses by having base-class pointers point to objects of derived classes. However, we may not write **programs** that are independent of the classes that will be used. Some place in the program, we must be explicit about exactly which classes are being used. These places in the program need to build a class and specify its initial values.

We would like to write whole programs that are independent of the explicit classes that will be used. If we do this we may add new classes to the code without touching code that has already been written. To do this we may extend the class to a component. A component is a run-time description of a class. A component knows how to make an object of the class, it knows

how to initialize the data members of the class, and it knows the base class of the class in order to determine the class's usage.

A component architecture has a single table where all such classes are registered. No other code needs to know specific classes that are used. Once linked in, the class is available and may be used anywhere its base class is used. Algorithms and functions need only refer to base classes and never need to know which objects are being used to represent these base classes.

Because data initialization requires knowledge of specific objects, and we would like to keep our programs independent of the specific objects that are used, we must find an alternative location to initialize objects other than in program source code. We use an input file to specify which components get built for Tecolote programs. Specifics of object initialization are determined at run time, rather than at program build time.

A Tecolote component is a C++ class with the special features MetaType and PERSISTENTS added.

The MetaType is the construct used to register a component with Tecolote. Once a MetaType is supplied for a class, that class is recognized as a component in Tecolote and new objects may be created from that component on the input file.

PERSISTENTS are used to initialize data members of a component from the input file. A class must specify which of its data members will be persistent and therefore accessible for initialization from the input file.

## Building a Tecolote Component

Building a component in Tecolote is similar to writing any other C++ class with a few notable exceptions. Our base-class shape need not be rewritten, but it needs to be specified as a base class in Tecolote. To do this one needs to create a MetaType for class shape.

```
static MetaClass<AbstractTraits<shape,
shape> >
shapeMeta("shape");
```

Class circle is similar to the class that was defined before with the following exception and additions. In the third step of creating a special method for constructing the class, we now need to construct an empty class with data members initialized to default values and declare that the class has persistent members. In the fifth step we specify which data members are persistent. And finally in the sixth step we build the MetaType for the class.

```
// 1. Make a class circle that inherits
from base-class shape
class circle : public shape {
// 2. Define data members of class
circle
float    x_center, y_center;
float    radius;
public:
// 3. Define Tecolote required members
PERSISTENT_MEMBERS(circle);
circle() { }
// 4. Define methods (functions) of
class circle

float area() { return 3.14 * radius *
radius; }

bool is_inside(float x_pt, float y_pt) {
```

```
return radius * radius <
pow(x_pt - x_center, 2 ) + pow(y_pt -
y_center 2 );
}
};

// 5. Define Persistents of class circle
BEGIN_PERSISTENTS(circle)
PERSISTENT(float, x_center,
"x_center")
PERSISTENT(float, y_center,
"y_center")
PERSISTENT(float, radius, "radius")
END_PERSISTENTS

// 6. Define the MetaType
static MetaClass<DefaultTraits<circle,
shape> >
circleMeta("circle",
MAKE_PERSISTENTS(circle));
```

A box may be similarly defined with the additional features to become a component.

This is a lot of overhead for a small class like circle, but for larger classes, this is a small percentage of the code needed.

MetaTypes may also be defined for functions. By defining a function as a MetaType, that function now becomes accessible from the input file. The exact parameters needed by the function will be deduced and a run-time check will be made to make sure the proper parameters are used when invoking the function. The following code shows an example function MetaType.

```
struct intersection_pointTag{ };

static MetaType*
intersection_pointMeta(makeFunction4(
intersection_pointTag(),
"intersection_point",
intersection_point));
```

## Using Tecolote Components to Initialize Objects at Run Time

As stated earlier, components are initialized from an input file that is processed at run time. The code below shows how two objects of class circle can be created from an input file.

```
c1 = circle{
x_center = 0.0
y_center = 0.0
radius   = 5.0
}
c2 = circle{
x_center = 1.0
y_center = 1.0
radius   = 2.0
}
```

The following input file creates two shapes and has a main function that checks if the point (1.0, 1.0) is in the intersection of the two shapes. The main function will be executed after the input file has been processed.

```
c = circle{
x_center = 1.0
y_center = 1.0
radius   = 0.5
}
```

```
b = box{
x_center = 0.0
y_center = 0.0
x_size  = 1.0
y_size  = 1.0
}

main = fcn is (Print intersection_point
1.0 1.0 c b)
```

### New Capabilities through Extending Tecolote Components

To date, Tecolote components are used to build an initial object hierarchy for a particular run of an application. In addition to this we would like to specify high-level control flow of a particular run of an application. To do this we are planning on using functional programming as a model for our scripting language with programming extensions. In support of this our components will be extended with methods in addition to having persistents. Because of these additions, class methods (member functions) will be accessible to the input file user in addition to class persistents (member data).

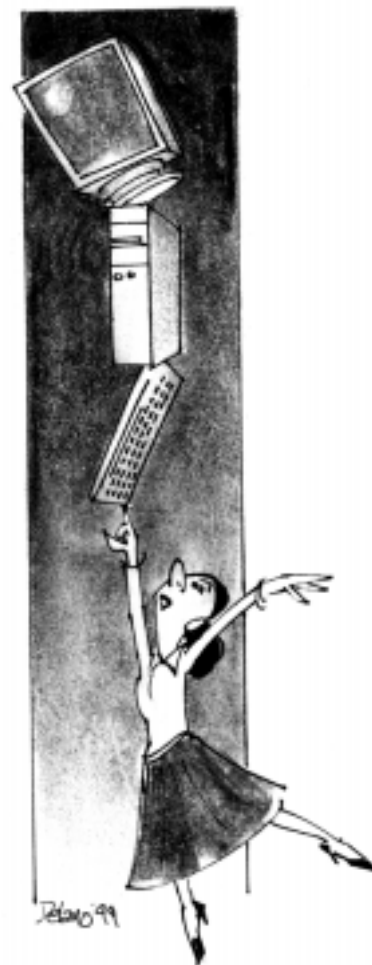
So far we've discussed using the class persistents to initialize objects from an input file. This same mechanism may be used for general purpose object I/O for Tecolote applications. For example, restart dumps may be written by traversing the object hierarchy and writing out the pertinent information. Currently we do this only to a limited extent but we would like to generalize the process. Since Tecolote components list the object attributes available for I/O through persistents, an I/O module may scan the object hierarchy writing out relevant information having no other knowledge of the application classes. Because of this I/O services may be completely independent of the applications that they serve.

There are additional services we would like to make available through components. The high-level synchronization needed for task parallelism could be represented by Tecolote components. Dynamic link libraries (DLLs) of components could be used to extend applications while they are running. These are just some of the areas where we would like to apply components in the future.

The Tecolote component architecture supports flexible applications by deferring the specification of object instantiation to run time through the use of an input file. We have found this to be advantageous over the previous inflexibility of programs with fixed implementations. Applications are viewed simply as collections of components. We have taken advantage of the flexible architecture by recombining components into new applications and experimenting with new components in existing applications.

The author wishes to acknowledge the project team: L.A. Ankeny, S.P. Clancy, W.H. Dorin, J.H. Hall, K.S. Holian, S.R. Lee, J.C. Marshall, G.R. McNamara, J.W. Painter, and S.J. Sydorik.

<sup>1</sup> For related article see *Tecolote: An Object-oriented Framework for Hydrodynamics Physics*, BITS August 1999 issue, at this URL: <http://www.lanl.gov/orgs/cic/computingatlanl/>.



# What's Happening

## CIC Division Hosts Student Poster Presentations

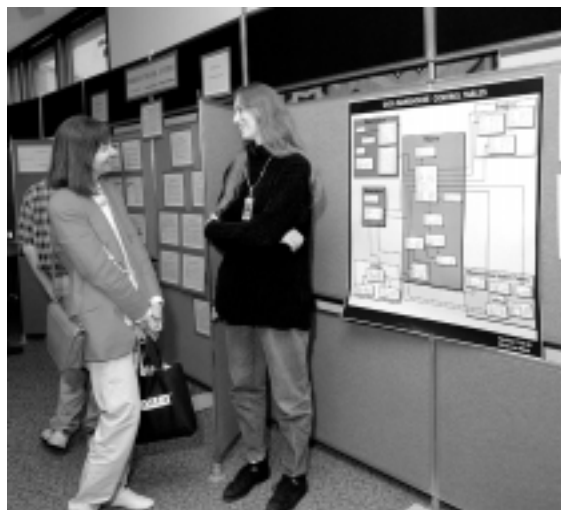
On the afternoon of August 6, 1999, CIC Division students presented poster sessions on their projects during their tenure at the Laboratory. According to the event's organizer, Kathy Hirons of CIC Division Office, "the Division wants to give students an opportunity to get recognition for their work." An annual event, this year's presentations included seventeen students representing thirteen poster sessions. In the following sampling of photographs taken during the event, note that not all students who worked on the projects are shown.



Towards Effluent Utilization of I/O Resources for Clusters of SMPs. *Karen Reid, Advanced Computing Laboratory.*



Computer Support at LANL Research Library. *Michelle Mirabal, Pete Zugger, CIC-14, Research Library (with Steve Peralta, Faculty Advisor, Santa Fe Community College).*



Data Warehouse Control Tables. *Rhiannon West, CIC-13, Business Information Systems.*



Foreign Travel System Components and Object Orientation. *Michael Bennett and Peter James Dennedy-Frank, CIC-15, Advanced Database and Information Technology.*



A Web Tool to Retrieve and Analyze Physics Data. *Paul Bailey, Mario Khalsa, Javier Olivares, and Greg Scudder, CIC-8, Distributed Computing.*



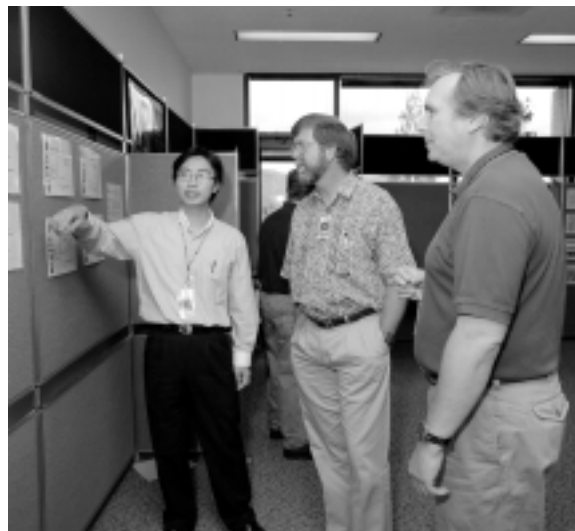
*Note: two poster presentations are shown in this photo. On the left: Scientific Visualization. David A. Hite, CIC-8, Distributed Computing. On the right: Preserving Legacy Data. Mark Bretten, CIC-8.*



*A Study in Java and RMI Technology. Mike Springer, CIC-8, Distributed Computing.*



*A Preliminary Study on the Impact of Memory Hierarchy on Stencil Operations in Modern Architectures. Federico Bassetti, CIC-19, Scientific Computing (with Kei Davis, Olaf Lubeck, Madham Marathe, Farbrizio Petrini, and Monica Reggiani).*



*Scal-Tools Pinpointing Scalability Bottlenecks. Yan Solihin, CIC-19, Scientific Computing Group (with Vinh Lam, Josep Torrellas, Yong Luo).*

## New Electronic Resources from the Research Library

### New INSPEC® at LANL

by Kathy Varjabedian, Database Librarian, CIC-14, Research Library

A new interface and search engine is now available for INSPEC® at LANL, our version of the world's leading database for physics, electronics, and computing. Features include:

- Alerts (current awareness service): weekly e-mail delivery of new information on subjects of your choice
- Easy limiting by date, language, document type, and electronic format
- Phrase and adjacency searching
- Sorting of results by relevancy ranking, date, author, or title
- Ability to search individual author names with initials
- Ability to mark "All" for printing, e-mail, or downloading
- Links to both printed and electronic Research Library holdings

The new system uses Verity's Topic search engine and looks much like SciSearch® at LANL. There are more indexes, because the records in INSPEC are much richer, containing detailed subject descriptors and subject classification codes. Also, whereas SciSearch contains only journal articles, INSPEC indexes journals, conferences, books, dissertations, and reports.

The INSPEC database has 6.3 million records covering the literature from 1969 to the present and is updated weekly. INSPEC scans papers from approximately 4,200 journals, 1000 conferences, and other publications, adding approximately 300,000 records a year. INSPEC corresponds to the printed publications *Physics Abstracts*, *Electrical and Electronics Abstracts*, and *Computer and Control Abstracts*.

---

### Keep on Top of Information Technology with Gartner Group

by Frances Knudson, Information Specialist, Library Without Walls Project, CIC-14, Research Library

GartnerGroup Interactive at <http://www.gartner.com> provides cutting-edge insights that support competitive decision making across the IT spectrum. Services and resources include comprehensive research, analysis, and intelligence; IT Journal, IT Vendor Directory, and information on hot topics such as software assessment management and business technology. The Web product allows one to search through the GartnerGroup content, to browse through selected publications, and to set up alerts on "My Homepage".

The Laboratory's contract with GartnerGroup has expanded to include all of the content in Gartner Advisory. Gartner Advisory contains the following categories:

- Administrative Applications Strategies
- Applications Development & Management Strategies
- Application Interface Middleware Strategies
- Advanced Technologies and

Applications

- Equipment Asset Management
- End-User Computing
- Integrated Document & Output Management
- Intranets & Electronic Workplace
- Internet Strategies
- Information Security Strategies
- Industry Trends & Directions
- Local Area Networking
- Managing Distributed Computing
- Management Strategies & Directions
- Network Computing Client/Server
- Networked Systems Management
- NT Strategies
- Rapid Development Solutions
- Return on IT Investment
- Software Data Management
- Storage Technologies, Operations, & Resources
- Telecommuting and Remote Access
- Year 2000 Strategies

More details on GartnerGroup Interactive and the content of the different advisory services and publications are available at <http://lib-www.lanl.gov/edata/ggdetails.htm>.

To access GartnerGroup, you will need the organization ID and password. These are available on the Library's sitemap <http://lib-www.lanl.gov/sitemap>.



# LOS ALAMOS NATIONAL LABORATORY Research Library

<http://lib-www.lanl.gov>

The LANL Research Library offers a variety of training opportunities for the Laboratory community. Available sessions focus on specialized library databases and other electronic resources. A complete list of course offerings can be found at <http://lib-www.lanl.gov/libinfo/training.htm>. All sessions are available to individuals or groups, at the library or your site. Arrange for a session by contacting the Library, phone 7-4175 or e-mail [library@lanl.gov](mailto:library@lanl.gov). Library tours are available on a drop-in basis every Wednesday at 1:00 p.m.

---

## Attention Application Developers!!!

**Need User Guides?**

**Need Online Help?**

CIC-1's Software Documentation Team can help you improve your product by developing user manuals, online help, quick reference cards, or other documentation to help your users learn the application. Contact Sheila Molony, 5-1585, or [sheila.molony@lanl.gov](mailto:sheila.molony@lanl.gov), for more information.



## Computer Training

The Customer Service Group (CIC-6) offers technical computer training (enterprise information applications, communications, office administration, and Web authoring) and advanced computer training (programming languages, system administration, and advanced applications). To register for a course access our Web page at <http://www.lanl.gov/internal/training/training.html>. Or from the LANL home page select the links: Training, Computer. For further information about technical computer training call (505) 667-9559, and for advanced technical computer training call (505) 667-9399.

Technical and Advanced Technical Computer Training Courses		
Communications	Office Skills 2000	Web Authoring and Browsing
<ul style="list-style-type: none"> <li>• Eudora 4.x</li> <li>• Meeting Maker 5.0.3</li> </ul>	<ul style="list-style-type: none"> <li>• Office Skills 2000–LANL</li> <li>• Office Skills 2000–Professional Development</li> </ul>	<ul style="list-style-type: none"> <li>• Dreamweaver 2.0–MAC or PC</li> <li>• FrontPage 2000</li> <li>• HTML Basics</li> <li>• HTML Intermediate</li> </ul>
Enterprise Information Applications (EIA)	Other EIA Courses	System Administration Training
<ul style="list-style-type: none"> <li>• Date Warehouse–Basics</li> <li>• Date Warehouse–EDS Reports</li> <li>• Data Warehouse–Passport</li> <li>• EDS —Basics</li> <li>• EDS–GUI</li> <li>• EDS–Training Plans</li> <li>• Infomaker</li> <li>• Invoice Approval System</li> <li>• Key/Core System</li> <li>• Purchase Card System</li> <li>• Procurement Desktop</li> <li>• Recharge</li> <li>• Time &amp; Effort GUI</li> <li>• Travel Foreign GUI</li> <li>• Travel Domestic GUI</li> <li>• Web JIT</li> </ul>	<ul style="list-style-type: none"> <li>• Financial Management Information System (FMIS)</li> <li>• Property Accounting, Inventory, and Reporting System (PAIRS)</li> <li>• Signature Authority System (SAS)</li> <li>• Secretarial/Contract Service (SE)</li> <li>• Salary Review System (SRS)</li> <li>• Directory Information System (DIS)</li> <li>• Automated Chemical Information System</li> </ul>	<ul style="list-style-type: none"> <li>• IRIX (SGI) System Administration (Beginning)</li> <li>• IRIX (SGI) System Administration (Advanced)</li> <li>• IRIX (SGI) Network Administration</li> <li>• SGI Performance Evaluation and System Tuning</li> <li>• Solaris 7 System Administration</li> <li>• Solaris 7 Network Administration</li> <li>• Solaris 7 Server Administration</li> <li>• Unix and Widows NT Integration</li> <li>• Windows NT Workstation and Server</li> <li>• Windows NT Optimization and Troubleshooting</li> <li>• Windows NT Security</li> </ul>
Programming Training	Application Training	ASCI
<ul style="list-style-type: none"> <li>• C Programming (Beginning)</li> <li>• C Programming (Advanced)</li> <li>• C++ for Experience C Programmers</li> <li>• ANSI/ISO C++ Programming Clinic (Advanced C++)</li> <li>• Distributed Objects Using Corba</li> <li>• Java Programming</li> <li>• Java Program Workshop</li> <li>• Distributed Programming with Java</li> <li>• Programming for Beginners Using Java</li> <li>• Object Technology–A Management Overview</li> <li>• Object-Oriented Analysis and Design</li> <li>• Perl Programming</li> <li>• Advanced Perl Programming with CGI</li> <li>• C-Shell Programming</li> <li>• Python for Scientists and Engineers</li> </ul>	<ul style="list-style-type: none"> <li>• Advanced WWW Development</li> <li>• FrameMaker Basic and Advanced</li> <li>• Foundations of IDL Programming</li> <li>• IDL 5.0 Graphic Object Workshop</li> <li>• Mastering MFC Development Using Visual C++ 6</li> <li>• Netscape Servers for Intranet Development</li> <li>• Origin 2000 Applications Programming and Optimization</li> <li>• Sendmail–Managing Internet Mail</li> <li>• C++ and the Unified Modeling Language</li> <li>• Sybase Fast Track to Adaptive Server Enterprise 11.5 (ASE)</li> <li>• Sybase Performance and Tuning for System 11</li> <li>• Sybase SQL Server Administration</li> <li>• Unix (Beginning)</li> <li>• Unix (Advanced)</li> <li>• Visual Basic 6.0 Fundamentals</li> <li>• Visual C++ Windows Programming</li> <li>• Windows NT Security</li> <li>• Windows 2000 Training</li> </ul>	<ul style="list-style-type: none"> <li>• Getting Started on ASCI Blue Mountain Systems</li> <li>• Running MPI on Blue Mountain Systems</li> <li>• Introduction to Totalview</li> <li>• LSF (Load Sharing Facility)</li> <li>• Introduction to HPSS (High Performance Storage System)</li> </ul>
* You do not need an ICN password to use e-mail.		

## Reader Feedback

Feedback helps us to provide a document that responds to the changing needs of its readership. If you have comments or questions about this publication, please let us hear from you. We have reserved the back of this form for that purpose. We also accept articles for publication that are of interest to our readers. Contact the managing editor for more information. This form is also used for new subscriptions, deletions, or changes. Instructions are on the back. If you prefer to contact us by e-mail, send your comments and/or subscription request to [denise@lanl.gov](mailto:denise@lanl.gov).

Do Not Staple  
Fold on This Line First



NO POSTAGE  
NECESSARY  
IF MAILED  
IN THE  
UNITED STATES

### BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 88 LOS ALAMOS NM

POSTAGE WILL BE PAID BY THE ADDRESSEE

MAIL STOP B251  
ATTN: DENISE SESSIONS, MANAGING EDITOR  
CUSTOMER SERVICE GROUP (CIC-6)  
LOS ALAMOS NATIONAL LABORATORY  
PO BOX 1663  
LOS ALAMOS, NM 87545-9916



Do Not Staple, Seal with Tape  
Fold Here

cut along dashed line

## Feedback

This image shows a single sheet of white paper with horizontal blue or grey ruling lines. The lines are evenly spaced and run across the width of the page. There are approximately 20 lines visible. The paper has a slight shadow on its right side, suggesting it's resting on a surface.

## New Subscription, Deletions, and Changes

Bits is published by Los Alamos National Laboratory. If you would like to be added to or deleted from our mailing list, please check the appropriate line, complete the form below, and mail us the form, or e-mail [bitsupdate@lanl.gov](mailto:bitsupdate@lanl.gov)

\_\_\_\_\_ Add my name to the BITS mail list.

\_\_\_\_\_ Delete my name from the BITS mailing list.

\_\_\_\_\_ Change my name/address as indicated below.

Name	Date	
Address	Mail Stop	
Group	Organization	
City	State	Zip
Phone	Number of copies	

## 1999 12-Month Index

For a more complete index, see [http://www.lanl.gov/orgs/cic/cic6/bits/index/index\\_home.html](http://www.lanl.gov/orgs/cic/cic6/bits/index/index_home.html)

Keywords	Title of BITS Article	Date	Page
ASCI	<i>Blue Mountain Is World's Fastest Computer</i>	Dec. 98	1
	<i>Tecolote: An Object-Oriented Framework for Hydrodynamics Physics</i>	Aug. 99	26
ASCI Training	<i>ASCI Training Course Descriptions</i>	J/J 99	27
BITS Announcements	<i>BITS Gets New Design, New Editor</i>	F/M 99	3
BITS Interviews	<i>Bob Newell</i>	O/N 98	1
	<i>Earleen Eden</i>	O/N 98	3
	<i>Dale Land</i>	Dec. 98	11
	<i>John Morrison</i>	F/M 99	6
	<i>Ann Hayes</i>	J/J 99	22
CIC (Computing, Information, & Communications)	<i>BITS Welcomes New Deputy Division Director for Strategic Computing</i>	F/M 99	6
CIC-2	<i>CIC-2's Future Is in Stewardship, Automation, and "Fun Stuff"</i>	Dec. 98	11
CIC-10	<i>CIC-10 "Knowledge Is Our Most Important Product"</i>	O/N 98	3
Desktop Computing	<i>CIC-2's Future Is in Stewardship, Automation, and "Fun Stuff"</i>	Dec. 98	11
	<i>New Ways for Managing Software Development Projects</i>	J/J 99	2
	<i>Scientific Software Process Engineering</i>	J/J 99	4
	<i>Adjusting for Reality: Mitigating Uncertainty in Projects</i>	Aug. 99	7
	<i>Your Data Is Gone . . . Now What?</i>	Aug. 99	11
Embedded Systems	<i>The Year 2000 Bug Is Hiding</i>	F/M 99	1
ESD	<i>Working for and Serving You</i>	F/M 99	10
HTML	<i>Extensible HTML</i>	J/J 99	12
High Performance Computing	<i>Performance and Scalability Analysis of Apps. on T-flop-scale Architectures</i>	J/J 99	7
	<i>Major 3-D Parallel Simulations</i>	J/J 99	9
	<i>Lecture Review: Next-Generation Chips, Processors, Transistors, and Wiring</i>	Aug. 99	23
	<i>Tecolote: An Object-Oriented Framework for Hydrodynamics Physics</i>	Aug. 99	26
Infrastructure	<i>Virtual Private Network</i>	J/J 99	15
	<i>Electronic Information Protection Regimes</i>	Aug. 99	13
	<i>Web Content Accessibility: New W3C Guidelines Have Arrived</i>	Aug. 99	20

Keywords	Title of BITS Article	Date	Page
Information Systems	<i>Engineering Index Is Now Available from the Research Library</i>	<i>Dec. 98</i>	<i>13</i>
	<i>Research Library Adds Social SciSearch® Database</i>	<i>F/M 99</i>	<i>5</i>
	<i>Library Databases</i>	<i>J/J 99</i>	<i>20</i>
	<i>aha!—An Interface to the Best of LANL Science &amp; Technology Web Resources</i>	<i>Aug. 99</i>	<i>2</i>
	<i>More Than a New Look &amp; Feel: the New LANL Badging System</i>	<i>Aug. 99</i>	<i>5</i>
Networks	<i>Laboratory Unclassified Network Will Implement Changes for Security</i>	<i>Dec. 98</i>	<i>3</i>
	<i>Preserving Bandwidth on Your Windows NT Network</i>	<i>F/M 99</i>	<i>8</i>
PNG	<i>It's Time for PNG: A Graphics Format You Can Pronounce</i>	<i>O/N 98</i>	<i>7</i>
Popular Science Award	<i>PS Awards LANL Web</i>	<i>O/N 98</i>	<i>12</i>
Records Management	<i>CIC-10: "Knowledge Is Our Most Important Product"</i>	<i>O/N 98</i>	<i>3</i>
SciSearch	<i>Research Library Adds Social SciSearch® Database</i>	<i>F/M 99</i>	<i>5</i>
Supercomputing	<i>Blue Mountain Is World's Fastest Computer</i>	<i>Dec. 98</i>	<i>1</i>
Virtual Private Network	<i>Virtual Private Network</i>	<i>J/J 99</i>	<i>15</i>
Wildfires	<i>Forecasting Wildfires and Other Crises</i>	<i>Apr. 98</i>	<i>3</i>
Windows NT	<i>Preserving Bandwidth on Your Windows NT Network</i>	<i>F/M 99</i>	<i>8</i>
World Wide Web (WWW or Web)	<i>It's Time for PNG: A Graphics Format You Can Pronounce</i>	<i>O/N 98</i>	<i>7</i>
	<i>Lab Web Publishers Invited to IntraLab99</i>	<i>F/M 99</i>	<i>7</i>
	<i>Extensible HTML</i>	<i>J/J 99</i>	<i>12</i>
Writing Resources	<i>Telling Your Story in BITS: Sharing Your Expertise and Energy</i>	<i>Aug. 99</i>	<i>31</i>
Year 2000 (Y2K)	<i>The Year 2000 Bug Is Hiding</i>	<i>F/M 99</i>	<i>1</i>

Customer Support Center . . . (505) 665-4444, ext. 851, or [cichelp@lanl.gov](mailto:cichelp@lanl.gov)

Because of a wide variety of CIC computing services, numerous facilities are available to address your questions. If you are uncertain whom to call, you can always call the Customer Support Center (CSC). CSC consultants are trained to either answer your question or locate someone who can. To reach the appropriate consultant, dial 665-4444 and make your selection from the following choices:

- Option 1: New user topics including e-mail, passwords, registration, and World Wide Web
- Option 2: Enterprise Information Applications such as Travel, Time and Effort, and Purchase cards
- Option 3: Scientific computing, storage systems, and networking
- Option 4: Classroom instruction and training
- Option 5: Desktop Consulting for PC and Macintosh software and network configurations.

### *Consulting Via E-Mail*

Customer Support .....	<a href="mailto:cichelp@lanl.gov">cichelp@lanl.gov</a>
Scientific and engineering computing .....	<a href="mailto:consult@lanl.gov">consult@lanl.gov</a>
Administrative and business computing ....	<a href="mailto:eiaconsult@lanl.gov">eiaconsult@lanl.gov</a>
Passwords and registration .....	<a href="mailto:validate@lanl.gov">validate@lanl.gov</a>
Macintosh computing .....	<a href="mailto:Mac-help@lanl.gov">Mac-help@lanl.gov</a>
PC computing .....	<a href="mailto:PC-help@lanl.gov">PC-help@lanl.gov</a>
UNIX computing .....	<a href="mailto:UNIX-help@lanl.gov">UNIX-help@lanl.gov</a>

### *Other Useful Numbers*

Advanced Computing Laboratory .....	665-4530
Central Computing Facility .....	667-4584
Network Operations Center .....	<a href="mailto:noc@lanl.gov">noc@lanl.gov</a> or 667-7423
Telephone Services Center .....	667-3400

*Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.*

*All company names, logos, and products mentioned herein are trademarks of their respective companies. Reference to any specific company or product is not to be construed as an endorsement of said company or product by The Regents of the University of California, the United States Government, The U.S. Department of Energy, nor any of their employees. The Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; therefore, the Laboratory as an institution does not endorse the viewpoint of a publication or guarantee its technical correctness.*

**Los Alamos**  
NATIONAL LABORATORY

# Los Alamos

NATIONAL LABORATORY

BITS is published monthly to highlight recent computing and communications activities within the Laboratory. We welcome your suggestions and contributions.

BITS may be accessed electronically at this URL:  
<http://www.lanl.gov/orgs/cic/computingatlanl/>

LALP-99-42 (9/99)